Massachusetts Institute of Technology
Department of Economics
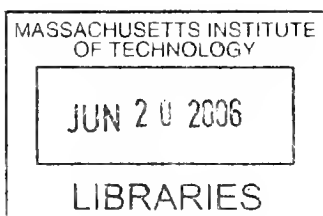Working Paper Series

CAPITAL TAXATION
Quantitative Exploration of the Inverse Euler Equation

Emmanuel Farhi
Iván Werning

Working Paper 06-15
November 30, 2005

Room E52-251
50 Memorial Drive
Cambridge, MA 02142

CAPITAL TAXATION[*]
Quantitative Explorations of the
Inverse Euler Equation

Emmanuel Farhi
MIT
efarhi@mit.edu

Iván Werning
MIT
iwerning@mit.edu

First Draft: November 2005
This Draft: May 2006

Preliminary and Incomplete
Comments Welcome

Abstract

This paper provides and employs a simple method for evaluating the quantitative importance of distorting savings in Mirrleesian private-information settings. Our exercise takes *any* baseline allocation for consumption—from US data or a calibrated equilibrium model using current policy—and solves for the best reform that ensures preserving incentive compatibility. The Inverse Euler equation holds at the new optimized allocation. Our method provides a simple way to compute the welfare gains and optimized allocation—indeed, yielding closed form solutions in some cases. When we apply it, we find that welfare gains may be quite significant in partial equilibrium, but that general equilibrium considerations mitigate the gains significantly. In particular, starting with the equilibrium allocation from Aiyagari's incomplete market model yields small welfare gains.

1

Introduction

Recent work has upset a cornerstone result in optimal tax theory. According to Ramsey models, capital income should eventually go untaxed (Chamley, 1986; Judd, 1985). In other words, individuals should be allowed to save freely and without distortions at the social rate of return to capital. This important benchmark has dominated formal thinking on this issue.

By contrast, in economies with idiosyncratic uncertainty and private information it is generally suboptimal to allow individuals to save freely: constrained efficient allocations satisfy an *Inverse Euler equation*, instead of the agent's standard intertemporal Euler equation (Diamond and Mirrlees, 1977; Rogerson, 1985; Ligon, 1998). Recently, extensions of this result have been interpreted as counterarguments to the Chamley-Judd no-distortion benchmark (Golosov, Kocherlakota and Tsyvinski, 2003; Albanesi and Sleet, 2004; Kocherlakota, 2004; Werning, 2002).

This paper explores the quantitative importance of these arguments. We examine how large welfare gains are from distorting savings and moving away from letting individuals save freely. In the process, we also extend and develop some new theoretical ideas. The issue we address is largely unexplored because—deriving first-order conditions aside—it is difficult to solve dynamic economies with private information, except for some very particular cases.[1]

Our approach sidesteps these difficulties by forgoing a complete solution for both consumption and work effort, and focusing, instead, entirely on consumption. Our exercise takes a given some baseline allocation—say, that generated in equilibrium with the current U.S. tax code—and improves on it in a way that does not affect the work effort allocation. That is, we optimize within a set of perturbations for consumption that guarantee preserving the incentive compatibility of the baseline work effort allocation. Our perturbations are rich enough to yield the Inverse Euler equation as a necessary condition for optimality, so that we fully capture recent arguments for distorting savings.

Our *partial reform* strategy is the relevant planning problem to address the capital taxation issue discussed above. That is, our calculations represent the welfare gains of moving away from allowing agents to save freely to a situation where their savings are optimally distorted.

In addition to being the relevant subproblem for the question addressed in this paper, there are some important advantages to a partial reform approach. First, the analysis of this problem turns out to be very tractable and flexible. As a result, one is not forced towards overly simplified assumptions, such as the specification of uncertainty, in order to make progress. In our view, this flexibility is important for any quantitative work.

---

[1] Two special cases that have been extensively explored are unemployment and disability insurance.

Secondly, and most importantly, our exercise does not require fully specifying some components of the economy. In particular, the details of the work-effort side of preferences and technology—such as how elastic work effort is to changes in incentives, whether the problem is one of private information regarding skills or of moral hazard regarding effort, etcetera—are not inputs in the exercise. This is a crucial advantage since current empirical knowledge of these things is limited and controversial. Our reforms are the richest ones one can consider without having to take a stand on these modelling details. Thus, we are able to address the question of how valuable it is to distort savings without evaluating the gains from selecting correctly along the tradeoff between insurance and incentives.

This paper relates to several strands of literature. First, there is the optimal taxation literature based on models with private information (see Golosov, Tsyvinski and Werning, 2006, and the references therein). Papers in this literature usually solve for constrained efficient allocations subject to the assumed asymmetry of information. Second, following the seminal paper by Aiyagari (1994), there is a vast literature on incomplete-market Bewley economies within the context of the neoclassical growth model. These papers emphasize the role of consumers self-smoothing through the precautionary accumulation of risk-free assets. In most positive analyses, government policy is either ignored or else a simple transfer and tax system is included and calibrated to current policies. In some normative analyses, some reforms of the transfer system, such as the income tax or social security, are evaluated numerically (e.g. Conesa and Krueger, 2005). Our paper bridges the gap between the optimal-tax and incomplete-market literatures by evaluating the importance of the constrained-inefficiencies in the latter.

The notion of efficiency used in the present paper is often termed *constrained-efficiency*, because it imposes the incentive-compatibility constraints that arise from the assumed asymmetry of information. Within exogenously incomplete-market economies, a distinct notion of constrained-efficiency has emerged (see Geanakoplos and Polemarchakis, 1985). The idea is roughly whether, taking the available asset structure as given, individuals could change their trading positions in such a way that generates a Pareto improvement at the resulting market-clearing prices. This notion has been applied by Davila, Hong, Krusell and Rios-Rull (2005) to Aiyagari's (1994) setup. They show that the resulting competitive equilibrium is inefficient. In this paper we also apply our methodology to examine an efficiency property of the equilibrium in Aiyagari's (1994) model, but it should be noted that our notion of constrained efficiency, which is based on preserving incentive-compatibility, is very different.

We proceed as follows. Section 1 lays out the basic backbone of our model, including assumptions regarding preferences, technology and information. Section 2 describes the perturbations and defines the planning problem we study. Section 3 derives a Bellman equation

representation, which serves as a methodological pillar for the rest of our analysis. Section 4 studies a partial equilibrium example that can be solved in closed form. Section 5 presents and employs the methodology for the general equilibrium problem. Section 6 extends our methodology from the infinitely lived agent model to an overlapping generations economy.

1   The Mirrleesian Economy

We cast our model within a general Mirrleesian dynamic economy. We follow the specification in Golosov, Kocherlakota and Tsyvinski (2003) most closely. They extend previous arguments leading to an Inverse Euler equation (e.g. Diamond and Mirrlees, 1977) by allowing individuals' privately observed skills to evolve as a general stochastic process. We will also show that our analysis applies to an extensions that includes a very general form of moral hazard.

Preferences. A continuum of agents have the utility function

$$\sum_{t=0}^{\infty} \beta^t \mathbb{E}[U(c_t) - V(n_t; \theta_t)] \tag{1}$$

where U is the utility function from consumption and V is the disutility function from effective units of labor (hereafter: labor for short). We assume U is increasing, concave and continuously differentiable. Additive separability between consumption and leisure is a feature of preferences that we adopt because it is required for the arguments leading to the Inverse Euler equation.[2]

Idiosyncratic uncertainty is captured by an individual specific shock $\theta_t \in \Theta$ that evolves as a Markov process. These shocks may be interpreted as determining skills or productivity, therefore affecting the disutility of effective units of labor. We denote the history up to period t by $\theta^t \equiv (\theta_0, \theta_1, \dots, \theta_t)$. Uncertainty is assumed idiosyncratic, so that the $\{\theta_t\}$ is independent and identically distributed across agents.

Allocations over consumption and labor are sequences $\{c_t(\theta^t)\}$ and $\{n_t(\theta^t)\}$. As is standard, it is convenient to change variables, translating any consumption allocation into utility assignments $\{u_t(\theta^t)\}$, where $u_t(\theta^t) \equiv U(c_t(\theta^t))$. This change of variable makes the incentive constraints below linear, rendering the planning problem convex.

Information and Incentives. The shock realizations are private information to the agent, so we must ensure that allocations are incentive compatible. We invoke the revelation prin-

---

[2] The intertemporal additive separability of consumption also plays a role. However, the intertemporal additive separability of work effort is completely immaterial: we could (pedantically) replace $\sum_{t=0}^{\infty} \beta^t \mathbb{E}[V(n_t; \theta_t)]$ with some general disutility function $\tilde{V}(\{n_t\})$.

4

ciple to derive the incentive constraints by considering a direct mechanism.

Imagine agents reporting each period their current shock realization. They are then allocated consumption and labor as a function of the entire history of reports. The agent's strategy determines a report for each period as a function of the history, $\{\sigma_t(\theta^t)\}$. The incentive compatibility constraint requires that truth-telling, $\sigma_t^*(\theta^t) = \theta_t$, be optimal:

$$\sum_{t=0}^{\infty} \beta^t \mathbb{E}[U(c_t(\theta^t)) - V(n_t(\theta^t); \theta_t)] \geq \sum_{t=0}^{\infty} \beta^t \mathbb{E}[U(c_t(\sigma^t(\theta^t))) - V(n_t(\sigma^t(\theta^t)); \theta_t)] \qquad (2)$$

for all reporting strategies $\{\sigma_t\}$.

Technology. We allow for a general technology over capital and labor, that includes the neoclassical formulation as a special case. Let K_t, N_t and C_t represent aggregate capital, labor and consumption for period t, respectively. The resource constraints are then

$$G(K_{t+1}, K_t, N_t, C_t) \leq 0 \qquad t = 0, 1, \dots$$

with initial capital K_0 given. The function $G(K', K, N, C)$ is assumed to be concave and continuously differentiable, increasing in K and N and decreasing in K' and C.

This technology specification is very flexible. For the neoclassical growth model we simply set $G(K', K, N, C) = -K' + F(K, N) + (1 - \delta)K - C$, where $F(K, N)$ is a concave constant returns production function satisfying Inada conditions $F_K(0, N) = \infty$ and $F_K(\infty, N) = 0$. Other technologies, such as convex adjustment costs, can just as easily be incorporated in the specification of G.

An important simple case is linear technology: labor can be converted to output linearly, with productivity normalized to one, and there is some storage technology with safe gross rate of return q^{-1}. One can think of this as partial equilibrium, perhaps representing the situation of a small open economy facing a fixed net interest rate $r = q^{-1} - 1$, or as the most stylized of A-K growth models (Rebelo, 1991). This special case fits our general setup with $G(K', K, N, C) = N - qK' + K - C$. With a continuum of agents, one can summarize the use of resources by their expected discounted value, using q as the discount factor.

Log Utility. Our model and its analysis, allows for general utility functions $U(c)$. However, the case with logarithmic utility $U(c) = \log(c)$ is of special interest for two reasons. The first reason is that we obtain simple formulas for the welfare gains in this case. Although the simplifications are useful, they are not in any way essential for our analysis.

More importantly, logarithmic utility is a natural benchmark for our model. Recall that, from the outset, we adopted additive separability between consumption and labor since this assumption is needed for the arguments leading to the Inverse Euler equation. But given

this separability a logarithmic utility function for consumption is a standard and sensible choice. In growth models, logarithmic utility is required for balanced growth. Moreover, a unitary coefficient of relative risk aversion and elasticity of substitution is broadly consistent with the range of empirical evidence.

Moral Hazard. We can extend our framework to introduce moral hazard as follows. Each period the agent takes an unobservable action e_t suffering some disutility, separable from consumption. Effective labor, n_t, is then a function of the history of effort and shocks, so that $n_t = f(e^t, \theta^t)$. Moreover, we can allow the distribution of θ^t to depend on past effort and shocks. Such a hybrid model seems relevant for thinking about the evolution of individuals' earned labor income—purely Mirrleesian or purely moral hazard stories are not likely to approximate reality very well.

2 Planning Problem

The starting point for our exercise is any allocation for consumption and labor that is incentive compatible. Given this baseline we consider a class of variations or perturbations that preserve incentive compatibility. The perturbations we consider induce shifts in the consumption allocation that ensure effort is unchanged. These perturbation are rich enough to deliver the Inverse Euler equation as an optimality condition. As a result, we fully capture the characterization of optimality stressed by Golosov et al. (2003).

2.1 Perturbations for Partial Reform

Parallel Perturbations. For any period t and history θ^t a feasible perturbation, of any baseline allocation, is to decrease utility at this node by $\beta\Delta$ and compensate by increasing utility by Δ in the next period for all realizations of θ_{t+1}. Total lifetime utility is unchanged. Moreover, since only parallel shifts in utility are involved, incentive compatibility of the new allocation is preserved. We can represent the new allocation as $\tilde{u}_t(\theta^t) = u_t(\theta^t) - \beta\Delta$ and $\tilde{u}_{t+1}(\theta^{t+1}) = u_{t+1}(\theta^{t+1}) + \Delta$, for all θ_{t+1}.

This perturbation changes the allocation in periods t and $t+1$ after history θ^t only. The full set of variations generalizes this idea by allowing perturbations of this kind at all nodes:

$$\tilde{u}(\theta^t) \equiv u(\theta^t) + \Delta(\theta^{t-1}) - \beta\Delta(\theta^t)$$

for all sequences of $\{\Delta(\theta^t)\}$ such that $\tilde{u}(\theta^t) \in U(\mathbb{R}_+)$ and such that the limiting condition

$$\lim_{T \to \infty} \beta^{T+1} \mathbb{E}[\Delta(\sigma^{t+T}(\theta^{t+T}))] = 0$$

for all reporting strategies $\{\sigma_t\}$. The limiting condition rules out Ponzi-like schemes in utility,[3] so that from *any* strategy $\{\sigma_t\}$ remaining expected utility is only changed by a constant Δ_{-1}:

$$\sum_{t=0}^{\infty} \beta^t \mathbb{E}[\tilde{u}_t(\sigma^t(\theta^t))] = \sum_{t=0}^{\infty} \beta^t \mathbb{E}[u_t(\sigma^t(\theta^t))] + \Delta_{-1}. \tag{3}$$

It follows directly from equation (2) that if the baseline allocation $\{u_t\}$ is incentive compatible then the new allocation $\{\tilde{u}_t\}$ remains so.[4] Note that the value of the initial shifter Δ_{-1} determines the lifetime utility of the new allocation relative to its baseline. In particular, if $\Delta_{-1} = 0$ then all perturbations yield the same utility as the baseline. In particular for any fixed infinite history $\bar{\theta}^{\infty}$ equation (3) implies that (by substituting the deterministic strategy $\sigma_t(\theta^t) = \bar{\theta}_t$)

$$\sum_{t=0}^{\infty} \beta^t \tilde{u}_t(\bar{\theta}^t) = \sum_{t=0}^{\infty} \beta^t u_t(\bar{\theta}^t) + \Delta_{-1}, \tag{4}$$

so that ex-post realized utility is the same along all possible realizations for the shocks.[5]

Let $\Upsilon(\{u_t\}, \Delta_{-1})$ denote the set of utility allocations $\{\tilde{u}_t\}$ that can be generated by such perturbations starting from a baseline allocation $\{u_t\}$ for a given initial Δ_{-1}. This is a convex set.

Partial Reform Problem. The problem we wish to consider takes as given a baseline allocation $\{u_t\}$ and an initial Δ_{-1} and seeks the alternative allocation $\{\tilde{u}_t\} \in \Upsilon(\{u_t\}, \Delta_{-1})$ that maximizes expected utility subject to the technological constraints. We want to allow for the possibility that in the initial period individuals are heterogeneous with respect to their shock, θ_0. We capture this without burdening the notation by letting the expectations operator integrate over the initial distribution for θ_0 in the population, in addition to future shocks.

---

[3] Note that the limiting restrictions are trivially satisfied for all variations with finite horizon: sequences for $\{\Delta_t\}$ that are zero after some period T. This was the case in the discussion of a perturbation at a single node and its successors.

[4] It is worth noting that in the logarithmic utility benchmark case parallel additive shifts in utility represent proportional shifts in consumption.

[5] The converse is nearly true: by taking appropriate expectations of equation (4) one can deduce equation (3), but for a technical caveat involving the possibility of inverting the order of the expectations operator and the infinite sum (which is always possible in a version with finite horizon and Θ finite). This caveat is the only difference between equation (3) and equation (4).

Planning Problem

$$\max_{\{\tilde{u}_t\}, C_t, K_{t+1}, \Delta_{-1}} \sum_{t=0}^{\infty} \beta^t \mathbb{E}[\tilde{u}_t(\theta^t)]$$

subject to,

$$\{\tilde{u}_t\} \in \Upsilon(\{u_t\}, \Delta_{-1})$$

$$G(K_{t+1}, K_t, N_t, C_t) \geq 0 \qquad C_t = \mathbb{E}[c(\tilde{u}_t(\theta^t))] \qquad t = 0, 1, \ldots \qquad (5)$$

where the cost function $c \equiv u^{-1}$ represents the inverse of the utility function $U(c)$.

This problem maximizes expected discounted utility within the set of feasible perturbations subject to an aggregate resource constraint (5). Note that from equation (3) we know that the objective function equals Δ_{-1} plus the utility from the baseline allocation.

2.2 Inverse Euler equation

In this section we review briefly the Inverse Euler equation which is the optimality condition for our planning problem.

Proposition 1. *At an interior optimum for the partial reform problem*

$$\beta c'(\tilde{u}_t) = q_t \mathbb{E}_t[c'(\tilde{u}_{t+1})] \quad \Rightarrow \quad \frac{1}{u'(\tilde{c}_t)} = \frac{q_t}{\beta} \mathbb{E}_t\left[\frac{1}{u'(\tilde{c}_{t+1})}\right]. \qquad (6)$$

where $q_t \equiv R_t^{-1}$ and

$$R_t \equiv -\frac{G_{C,t}}{G_{K',t}} \frac{G_{K,t+1}}{G_{C,t+1}}$$

is the technological rate of return (i.e. the technologically feasible tradeoff for dC_{t+1}/dC_t).

Equation (6) is the so called *Inverse Euler equation*. The important point is that our perturbations are rich enough to recover the crucial optimality condition stressed in Golosov et al. (2003).

It is useful to contrast equation (6) with the standard Euler equation that obtains in incomplete market models. In models where agents can save freely at the rate of return to capital we obtain

$$u'(\tilde{c}_t) = \beta R_t \mathbb{E}_t[u'(\tilde{c}_{t+1})]. \qquad (7)$$

Note that equation (6) is a reciprocal of sorts version of equation (7). Applying Jensen's inequality to equation (6) implies that at the optimum

$$u'(\tilde{c}_t) < \beta R_t \, \mathbb{E}_t[u'(\tilde{c}_{t+1})]. \tag{8}$$

To use Rogerson's (1985) language: at the optimum equation (8) implies that individuals are 'savings constrained'. Alternatively, another interpretation is that at the optimum a version of equation (7) can hold if we substitute the social rate of return R_t for a lower private rate of return—this difference in such rates of return is has been termed 'distortion', 'wedge' or 'implicit tax'.

Regardless of the interpretation, the optimum cannot allow agents to save freely at technology's rate of return, since then equation (7) would hold as a necessary condition, which we know is incompatible with the planner's optimality condition, equation (6). This is in stark contrasts with the Chamley-Judd benchmark obtained in Ramsey models, where it is optimal to let agents save freely at the technological rate of return to capital.

We now consider a subproblem that takes as given aggregate consumption and savings, and only optimizes over the distribution of this aggregate consumption. Of course, this is the only relevant problem for version of the economy without capital; but it is also of interest because of its relation to some welfare decompositions we shall perform later. We recover a variant of equation (6) equivalent to that obtained by Ligon (1998).

Proposition 2. *Consider the subproblem that replaces the resource constraints in (5) with the constraint that $\mathbb{E}[c(\tilde{u}_t(\theta^t))] \le \hat{C}_t$ for $t = 0, 1, \ldots$, where $\{\hat{C}_t\}$ is any given sequence for aggregate consumption. Then at an interior optimum there must exist a sequence of discount factors $\{q_t\}$ such that for all histories θ^t:*

$$\mathbb{E}_t\left[\frac{u'(\tilde{c}_t)}{u'(\tilde{c}_{t+1})}\right] = \frac{\beta}{q_t} \qquad t = 0, 1, \ldots \tag{9}$$

Equation (9) simply states that the expected ratio of marginal utilities is equalized across agents and histories, but not necessarily linked to any technological rate of transformation. Conditions equalizing marginal rates of substitution are standard in exchange economies. However, unlike the standard condition obtained in an incomplete market equilibrium model, here it is the expectation of $u'(c_t)/u'(c_{t+1})$, instead of $u'(c_{t+1})/u'(c_t)$, that is equated. Note that the implication in equation (6) is stronger, since it implies equation (9) but also pins down q_t.

2.3 Steady States with CRRA

We now focus on steady states for the planning problem, where aggregates are constant over time. For purposes of comparison, we discuss the conditions in the context of the neoclassical growth model specification of technology. We also limit our discussion to constant relative risk aversion (CRRA) utility functions $U(c) = c^{1-\sigma}/(1-\sigma)$. Finally, we suppose that the baseline allocation features constant aggregate labor $N_t = \bar{N}$.

At a steady state level of capital K_{ss} the discount factor is constant $q_t = q_{ss} \equiv 1/(1+r_{ss})$ where

$$r_{ss} = F_K(K_{ss}, \bar{N}) - \delta.$$

With CRRA preferences the Inverse Euler equation is

$$\mathbb{E}_t[\tilde{c}_{t+1}^\sigma] = (1 + r_{ss})\beta \tilde{c}_t^\sigma \tag{10}$$

At a steady state we must have constant consumption $\tilde{C}_t = \tilde{C}_{ss}$. With logarithmic utility $\sigma = 1$ the equation above implies, by the law of iterated expectations, that

$$\tilde{C}_{t+1} \equiv \mathbb{E}[\tilde{c}_{t+1}] = (1 + r_{ss})\beta \mathbb{E}[\tilde{c}_t] \equiv (1 + r_{ss})\beta \, \tilde{C}_t.$$

Hence, consumption is constant if and only if $(1 + r_{ss})\beta = 1$. For $\sigma > 1$, we can use Jensen's inequality to obtain $\mathbb{E}_t[\tilde{c}_{t+1}] \leq (\beta/q)^{1/\sigma}\tilde{c}_t$, so it follows that $\tilde{C}_{t+1} \leq (\beta/q)^{1/\sigma}\tilde{C}_t$. Then, if $(1 + r_{ss})\beta < 1$ aggregate consumption $\tilde{C}_t$ would strictly fall over time, contradicting a steady state with constant positive consumption. The argument for $\sigma < 1$ is symmetric.

Proposition 3. *Suppose CRRA preferences $U(c) = c^{1-\sigma}/(1-\sigma)$ with $\sigma > 0$. Then at a steady state of the planning problem*

(i) if $\sigma \geq 1$ then $r_{ss} \geq \beta^{-1} - 1$

(ii) if $\sigma \leq 1$ then $r_{ss} \leq \beta^{-1} - 1$

We depict steady state equilibria by adapting a graphical device introduced by Aiyagari (1994). For any given interest rate r we solve a fictitious planning problem with a linear savings technology with $q \equiv (1+r)^{-1}$; equivalently, this is the partial equilibrium. Intuitively, for $q_{ss} = (1 + r_{ss})^{-1}$ the steady state for the fictitious planning coincides with the for the actual general equilibrium planning problem—Section 5.3 formalizes this notion.

To find the equilibrium we imagine solving these fictitious planning problems for all arbitrary r. Let $C(r)$ denote the average steady state consumption as a function of r; one can show that this function is increasing, which is intuitive since a higher r makes
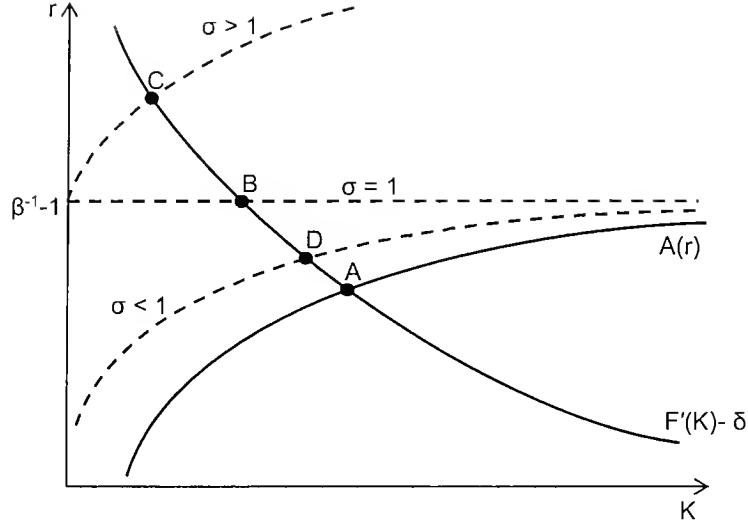
Figure 1: Steady States: Market Equilibrium vs. Planner. The intersection point A represents the baseline market equilibrium. Points B, C and D represent, respectively, the planner's steady states for log utility, $\sigma > 1$ and $\sigma < 1$.

providing future (long run) consumption cheaper to the planner. Define $K(r)$ to be the solution to the steady state resource constraint $C(r) = F(K, \bar{N}) - \delta k$, for $k \leq k_g$, where $k_g \equiv \arg\max_K \{F(K, \bar{N}) - \delta K\}$ is the golden rule level of capital.[6]

In Figure 1 the unique steady state capital and interest rate (K_{ss}, r_{ss}) is then found at the intersection of the downward sloping technological schedule $r = F_K(K, \bar{N}) - \delta$ with the upward sloping $K(r)$ schedule. For comparison, the figure also displays the equilibrium schedule for steady-state average savings, $A(r)$, that is obtained from the incomplete markets model in Aiyagari (1994). The steady state equilibrium level of capital is always higher than the steady state optimum; the interest rate r is always lower.

3   A Useful Bellman Equation

We now study the case with a constant rate of discount q. The resource constraints can then be reduced to a single expected present value condition. One literal interpretation is that this represents an economy with a linear saving technology. It also represents a 'partial-equilibrium' component of the 'general-equilibrium' planning problem. In any case, the analysis we develop here will be useful later in Section 5, when we return to the problem

---

[6]If r is so high that $C(r) > F(k_g, \bar{N}) - \delta k_g$ then $C(r)$ cannot be sustainable as a steady state. Hence, we limit ourselves to values of r for which $C(r) \leq F(k_g, \bar{N}) - \delta k_g$.

with general technology.

With a fixed discount factor q the (dual of the) planning problem minimizes the expected discounted cost

$$\sum_{t=0}^{\infty} q^t \mathbb{E}[c(\tilde{u}_t)] = \sum_{t=0}^{\infty} q^t \mathbb{E}[c(u_t + \Delta_t - \beta\Delta_{t+1})]$$

Let $K(\Delta_{-1}; \theta_0)$ denote the lowest achievable expected discounted cost over all feasible perturbation $\{\tilde{u}_t\} \in \Upsilon(\{u_t\}, \Delta_{-1})$. Since both the objective and the constraints are convex it follows immediately that the value function $K(\Delta_{-1}; \theta_0)$ is convex in Δ_{-1}.

Recursive Baseline Allocations. We wish to approach the planning problem recursively. This leads us to be interested in situations where the effects shocks θ^t on the baseline allocation can be summarized by some state s_t that evolves as a Markov process. In any period the assignment of utility is determined by this state and write the allocation as $u(s_t)$.

The requirement that the baseline allocation be expressible in terms of some state variable s_t is hardly restrictive. It is satisfied by virtually all baseline allocations of interest. To see this, recall that the exogenous shock θ_t is a Markov process. Most economic models then imply that consumption is a function of the full state $s_t = (x_t, \theta_t)$, where x_t is some endogenous state with law of motion $x_t = G(x_{t-1}, \theta_t)$. The endogenous state and its law of motion depends on the particular economic model generating the baseline allocation. A leading example in this paper is the case of incomplete market models, such as the Bewley economies in Huggett (1993) and Aiyagari (1994). In these models, described in more detail in Section 5.2, individuals are subject to exogenous shocks to income or productivity and can save using a riskless asset. The endogenous state for any individual is their asset wealth and the law of motion is their optimal saving rule.[7]

A Bellman equation. We now reformulate the above planning problem recursively for such recursive baseline allocations. First, since the problem can be indexed by s_0 instead of θ_0 we rewrite the value function as $K(\Delta_{-1}, s_0)$. The value function K must satisfy the following Bellman equation.

Bellman equation

$$K(\Delta; s) = \min_{\Delta'} \left[c(u(s) + \Delta - \beta\Delta') + q\mathbb{E}[K(\Delta'; s') \mid s] \right] \tag{11}$$

Moreover, the new optimized allocation from the sequential partial reform problem must be

---

[7] Another example are allocations generated by some dynamic contract. The full state variable then includes the promised continuation utility (see Spear and Srivastava, 1987) along with the exogenous state.

generated by the policy function from the Bellman equation.

Note that from the planner's point of view the state s_t evolves exogenously. Thus, the recursive formulation takes s_t as an exogenous shock process and employs Δ as an endogenous state variable. In words, the planner keeps track of the additional lifetime utility Δ_{t-1} it has previously promised the agent up and above the welfare entitled by the baseline allocation. Note that the resulting dynamic program is very simple: the endogenous state variable Δ is one-dimensional and the optimization is convex.

Inverse Euler equation (again). The first-order and envelope conditions are

$$\beta c'(u(s) + \Delta - \beta\Delta') = q\mathbb{E}[K_\Delta(\Delta'; s') \mid s],$$
$$K_\Delta(\Delta; s) = c'(u(s) + \Delta - \beta\Delta').$$

Combining these two conditions leads to the *Inverse Euler equation* equation (6).

A Useful Analogy. The planning problem admits an analogy with a consumer's income fluctuation problem that is both convenient and enlightening. We transform variables by changing signs and switch the minimization to a maximization. Let $\tilde{\Delta} \equiv -\Delta$, $\tilde{K}(\tilde{\Delta}; s) \equiv -K(-\tilde{\Delta}; s)$ and $\tilde{U}(x) \equiv -c(-x)$. Note that the pseudo utility function $\tilde{U}$ is increasing, concave and satisfies Inada conditions at the extremes of its domain.[8] Reexpressing the Bellman equation (11) using these transformation yields:

$$\tilde{K}(\tilde{\Delta}; s) = \max_{\tilde{\Delta}'} \left[\tilde{U}(-u(s) + \tilde{\Delta} - \beta\tilde{\Delta}') + q\mathbb{E}[\tilde{K}(\tilde{\Delta}'; s') \mid s] \right].$$

This reformulation can be read as the problem of a consumer with a constant discount factor q facing a constant gross interest rate $1 + r = \beta^{-1}$, entering the period with pseudo financial wealth $\tilde{\Delta}$, receives a pseudo labor income shock $-u(s)$. The fictitious consumer must decide how much to save $\beta\tilde{\Delta}'$; pseudo consumption is then $x = -u(s) + \tilde{\Delta} - \beta\tilde{\Delta}'$.

The benefit of this analogy is that the income fluctuations problem has been extensively studied and used; it is at the heart of most general equilibrium incomplete market models (Aiyagari, 1994). Intuition and results have accumulated with long experience.

Log Magic. With logarithmic utility, a case we argued earlier constitutes an important benchmark, the Bellman equation can be simplified considerably. The idea is best seen through the analogy, noting that the pseudo utility function is exponential $-e^{-x}$ in this

---

[8] An important case is when the original utility function is CRRA $U(c) = c^{1-\sigma}/(1-\sigma)$ for $\sigma > 0$ and $c \geq 0$. Then for $\sigma > 1$ the function $\tilde{U}(x)$ is proportional to a CRRA with coefficient of relative risk aversion $\tilde{\sigma} = \sigma/(\sigma - 1)$ and $x \in (0, \infty)$. For $\sigma < 1$ the pseudo utility $\tilde{U}$ is "quadratic-like", in that it is proportional to $-(-x)^\rho$ for some $\rho > 1$, and $x \in (-\infty, 0]$.

case. It is well known that for a consumer with CARA preferences a one unit increase in financial wealth, $\tilde{\Delta}$, results in an increase in pseudo-consumption, x, of $r/(1+r) = 1 - \beta$ in parallel across all periods and states of nature. It is not hard to see that this implies that the value function takes the form $\tilde{K}(\tilde{\Delta}; s) = e^{-(1-\beta)\tilde{\Delta}}\tilde{k}(s)$.

Of course, this decomposition translates directly to the original value function. Less obvious is that it greatly simplifies the Bellman equation and its solution.

Proposition 4. *With logarithmic utility the value function is given by*

$$K(\Delta; s) = e^{(1-\beta)\Delta}k(s),$$

where function $k(s)$ solves the Bellman equation

$$k(s) = Ac(s)^{1-\beta}\left(\mathbb{E}[k(s') \mid s]\right)^{\beta}, \tag{12}$$

where $A \equiv (q/\beta)^{\beta}/(1-\beta)^{1-\beta}$.

The optimal policy for Δ can be obtained from $k(s)$ using

$$\Delta' = \Delta - \frac{1}{\beta}\log\left(\frac{(1-\beta)k(s)}{c(s)}\right). \tag{13}$$

Proof. See Appendix A. ∎

This solution is nearly closed form: we need only compute $k(s)$ using the recursion in equation (12), which requires no optimization. It is noteworthy that no simplifications on the stochastic process for skills are required.

4 Example: Random Walk in Partial Equilibrium

Although the main virtue of our approach is that we can flexibly apply it to various baseline allocations (and we will!) in this section we begin with a simple and instructive case. We take the baseline allocation to be a geometric random walk and center our discussion around logarithmic utility. The advantage is that we obtain closed-form solutions for the optimized allocation, the intertemporal wedge and the welfare gains. The transparency of the exercise reveals important determinants for the magnitude of welfare gains.

Although extremely stylized, a random walk is an important conceptual and empirical benchmark. First, most theories—starting with the simplest permanent income hypothesis—predict that consumption should be close to a random walk. Second, some authors have

argued that the empirical evidence on income, which is a major determinant for consumption, and consumption itself shows the importance of a highly persistent component (e.g. Storesletten, Telmer and Yaron, 2004). For these reasons, a parsimonious statistical specification for consumption may favor a random walk.

An Example Economy. Indeed, one can construct an example economy where a geometric random walk for consumption arises as a competitive equilibrium with incomplete markets. To see this, suppose that individuals have CRRA utility over consumption—logarithmic utility being a special case—and disutility $V(n; \theta) = v(n/\theta)$ for some convex function $v(n)$ over work effort n, so that θ can be interpreted as productivity. Skills evolve as a geometric random walk, so that $\theta_{t+1} = \varepsilon_{t+1}\theta_t$, where ε_{t+1} is i.i.d. Individuals can only accumulate a riskless asset paying return R equal to the rate of return on the economy's linear savings technology. They face the sequence of budget constraints

$$a_{t+1} + c_t \leq Ra_t + \theta_t n_t \qquad t = 0, 1, \dots$$

and the borrowing constraint that $a_t \geq 0$. Suppose the rate of return R is such that $1 \geq \beta R \mathbb{E}[\varepsilon^{-1}]$. Finally, suppose that individuals have no initial assets, so that $a_0 = 0$.

In the competitive equilibrium of this example individuals exert a constant work effort $\bar{n}$, satisfying $\bar{n}v'(\bar{n}) = 1$, and consume all their labor income each period, $c_t = \theta_t \bar{n}$; no assets are accumulated, a_t remains at zero. This follows because the construction ensured that the agent's intertemporal Euler equation holds at the proposed equilibrium consumption process. The level of work effort $\bar{n}$ is defined so that the intra-period consumption-leisure optimality condition holds. Then, since the agent's problem is convex, it follows that this allocation is optimal for individuals. Since the resource constraint trivially holds, it is an equilibrium.

Although this is certainly a very special example economy, it illustrate that a geometric random walk for consumption is a possible equilibrium outcome.

Value Function. We now simply assume that at the baseline $s' = \varepsilon \cdot s$ with ε i.i.d. and $c = s$, so that $u(s) = U(s)$. The next result follows almost immediately from Proposition 4.

Proposition 5. *If consumption is a geometric random and the utility function is logarithmic, then the value function and policy functions are*

$$K(\Delta; s) = c(s)\frac{1}{1 - \beta}\exp((1 - \beta)\Delta)\, g^{\frac{\beta}{1-\beta}}, \tag{14}$$

$$\Delta' = \Delta - \frac{1}{1 - \beta}\log(g), \tag{15}$$

where $g \equiv (q/\beta)\mathbb{E}[\varepsilon]$.

Proof. With logarithmic utility the value function is of the form $K(\Delta, s) = \exp((1 - \beta)\Delta)k(s)$ where $k(s)$ given by Proposition 4. With a geometric random walk the problem is homogeneous and it follows that $k(s) = \kappa s$, for some constant κ. Solving for κ using equation (12), gives equation (14). Then Equation (13) delivers equation (15). ∎

In this example, the Inverse Euler equation is simply $g \equiv (q/\beta)\mathbb{E}[\varepsilon] = 1$. Hence, $g - 1$ measures the departure of the baseline allocation from it. We next show that $g - 1$ is critical in determining the magnitude of the intertemporal corrective wedge and the welfare gains from the optimized allocation.

Allocation. Using equation (15) we can express the new optimized consumption as

$$c(s, \Delta) = \exp(U(s) + \Delta - \beta\Delta') = g^{\frac{\beta}{1-\beta}} s \exp((1 - \beta)\Delta).$$

Since equation (15) also implies that

$$\exp((1 - \beta)\Delta') = g^{-1} \exp((1 - \beta)\Delta),$$

it follows that the new consumption allocation is also a geometric random walk. Indeed, it equals the baseline up to a multiplicative deterministic factor that grows exponentially:

$$\tilde{c}_t = \frac{\beta}{q} \frac{\varepsilon_t}{\mathbb{E}[\varepsilon_t]} \tilde{c}_{t-1} = \alpha g^{-t} c_t. \tag{16}$$

The drift in the new consumption allocation is β/q, ensuring that the Inverse Euler equation $\tilde{c}_t = (q/\beta)\mathbb{E}_t[\tilde{c}_{t+1}]$ holds at the optimized allocation.

Intertemporal Wedge. Suppose the agent's Euler equation holds at the baseline allocation:

$$1 = \frac{\beta}{q}\mathbb{E}\left[\frac{c_t}{c_{t+1}}\right] = \frac{\beta}{q}\mathbb{E}\left[\varepsilon^{-1}\right]. \tag{17}$$

We can then solve for the intertemporal wedge τ at the new allocation that makes the agent's euler equation hold:

$$1 = \frac{\beta}{q}(1 - \tau)\mathbb{E}\left[\frac{\tilde{c}_t}{\tilde{c}_{t+1}}\right] = \frac{\beta}{q}(1 - \tau)g\mathbb{E}\left[\varepsilon^{-1}\right]. \tag{18}$$

By using equation (16) in combination with equations (17) and (18) yield

$$\tau = 1 - \frac{1}{g}. \tag{19}$$

Applying Jensen's inequality to equation (17) gives

$$\frac{1}{g} = \frac{\beta}{q}\frac{1}{\mathbb{E}[\varepsilon]} < \frac{\beta}{q}\mathbb{E}\left[\varepsilon^{-1}\right] = 1 \quad \Rightarrow \quad \tau > 0.$$

This example the Inverse Euler equation provides a rationale for a constant and positive wedge in the agent's Euler equation. This is in stark contrast to the Chamley-Judd benchmark result, where no such distortion is optimal in the long run, so that agents are allowed to save freely at the social rate of return.

Welfare Gains. Evaluating the value function, given by equation (14), at $\Delta = 0$ gives the expected discounted cost of the new optimized allocation. The baseline allocation, on the other hand, costs $s\psi$ with

$$s\psi = s + q\mathbb{E}[s\varepsilon\psi] \quad \Rightarrow \quad \psi = \frac{1}{1 - q\mathbb{E}[\varepsilon]}.$$

The relative reduction in costs is then

$$\frac{s\psi}{sk(0)} = \frac{\psi}{k(0)} = \frac{1-\beta}{1-q\mathbb{E}[\varepsilon]}g^{-\frac{\beta}{1-\beta}} = \frac{\beta^{-1}-1}{\beta^{-1}-g}\ g^{-\frac{\beta}{1-\beta}}. \tag{20}$$

By homogeneity, the ratio of costs does not depend on the level of the current shock s. Note that at $g = 1$ there are no cost reductions. This is because, as discussed above, in this case the Inverse Euler holds at the baseline allocation. At the other extreme, as $g \to \beta^{-1}$ the cost reductions become arbitrarily large. The reason is that then $q\mathbb{E}[\varepsilon] \to 1$, implying that the present value of the baseline consumption allocation goes to infinity; in contrast, $k(0)$ remains finite.

The variance of consumption growth. Given the importance of g, we now investigate its main determinants. We can back out g if assume the agent's Euler equation holds at the baseline allocation and that ε is lognormally distributed. The latter assumption implies that

$$\mathbb{E}\left[\varepsilon^{-1}\right] = \exp\left(-\mu + \frac{\sigma_\varepsilon^2}{2}\right) \qquad \text{and} \qquad \mathbb{E}\left[\varepsilon\right] = \exp\left(\mu + \frac{\sigma_\varepsilon^2}{2}\right).$$

$$\Rightarrow \quad \mathbb{E}\left[\varepsilon\right] \cdot \mathbb{E}\left[\varepsilon^{-1}\right] = \exp(\sigma_\varepsilon^2).$$

Multiplying and dividing the agent's Euler equation, $(\beta/q)\mathbb{E}[\varepsilon^{-1}] = 1$, by $\mathbb{E}[\varepsilon]$ and rearranging:

$$g = \mathbb{E}\left[\varepsilon\right] \cdot \mathbb{E}\left[\varepsilon^{-1}\right] = \exp(\sigma_\varepsilon^2)$$
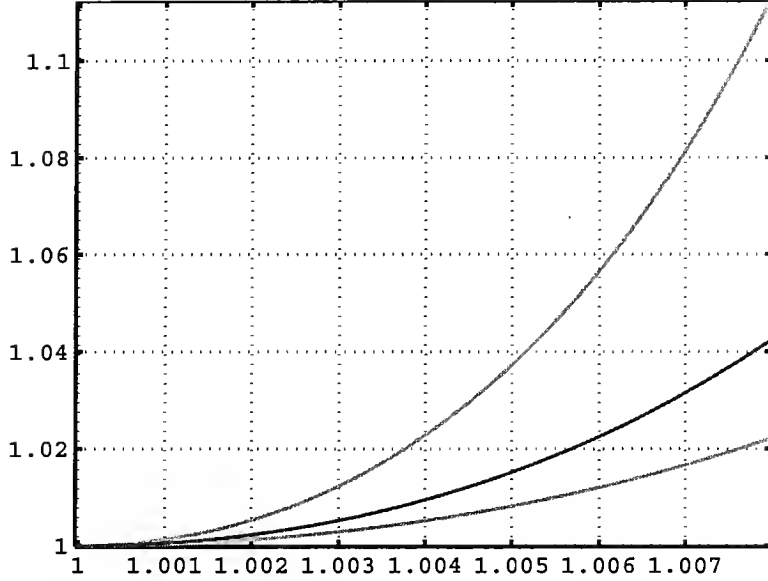$$\Rightarrow \quad g \approx 1 + \sigma_\varepsilon^2.$$

Figure 2: Welfare gains when baseline consumption is a geometric random walk for consumption as function of $g \approx 1 + \sigma_\varepsilon^2$—see equation (20). The blue line shows the case with $\beta = .97$, while the top red line has $\beta = .98$ and the bottom red line has $\beta = .96$.

Hence, we find that the crucial parameter g is associated with the variance in the growth rate of consumption. Equation (19) gives $\tau = 1 - \exp(-\sigma_\varepsilon^2)$, or approximately

$$\tau \approx \sigma_\varepsilon^2. \tag{21}$$

Thus, the intertemporal wedge τ is approximately equal to the variance in the growth rate of consumption σ_ε^2.

Quantitative Stab. Figure 2 plots the reciprocal of the relative cost reduction using equation (20), a measure of relative welfare gains. The blue line is for $\beta = .97$, while the top and bottom red lines are for $\beta = .98$ and $\beta = .96$, respectively. The figure uses an empirically relevant range for $g \approx 1 + \sigma_\varepsilon^2$.

In the figure the effect of the discount factor β is nearly equivalent to increasing the variance of shocks; that is, moving from $\beta = .96$ to $\beta = .98$ has the same effect as doubling σ_ε^2. To understand this, interpret the lower discounting not as a change in the actual subjective discount, but as calibrating the model to a shorter period length. But then holding the variance of the innovation between periods constant implies an increase in uncertainty over any fixed length of time. Clearly, what matters is the amount of uncertainty per unit of discounted time.

Empirical Evidence. Suppose one wishes to accept the random walk specification of con-

sumption as a useful empirical approximation. What does the available empirical evidence say about the crucial parameter σ_ε^2?

Unfortunately, the direct empirical evidence on the variance of consumption growth is very scarce, due to the unavailability of good quality panel data for broad categories of consumption.[9] Moreover, much of the variance of consumption growth in panel data may be measurement error or attributable to transitory taste shocks unrelated to the permanent changes we are interested in here.

There are a few papers that, somewhat tangentially, provide some direct evidence on the variance of consumption growth. We briefly review some of this recent work to provide a sense of what is currently available. Using PSID data, Storesletten, Telmer and Yaron (2004) find that the variance in the growth rate of the permanent component of food expenditure lies between 1%–4%.[10] Blundell, Pistaferri and Preston (2004) use PSID data, but impute total consumption from food expenditure. Their estimates imply a variance of consumption growth of around 1%.[11] Krueger and Perri (2004) use the panel element in the Consumer Expenditure survey to estimate a statistical model of consumption. At face value, their estimates imply enormous amounts of mobility and a very large variance of consumption growth—around 6%–7%—although most of this should be attributed to a transitory, not permanent, component.[12] In general, these studies reveal the enormous empirical challenges faced in understanding the statistical properties of household consumption dynamics from available panel data.

An interesting indirect source of information is the cohort study by Deaton and Paxson (1994). This paper finds that the cross-sectional inequality of consumption rises as the cohort ages. The rate of increase then provides indirect evidence for σ_ε^2; their point estimate implies a value of $\sigma_\varepsilon^2 = 0.0069$. However, recent work using a similar methodology finds much lower estimates (Slesnick and Ulker, 2004; Heathcote, Storesletten and Violante, 2004b).

We conclude that, in our view, the available empirical evidence does not provide reliable estimates for the variance of the permanent component of consumption growth. Thus, for our purposes, attempts to specify the baseline consumption process directly are impractical. That is, even if one were willing to assume the most stylized and parsimonious statistical

---

[9] For the United States the PSID provides panel data on food expenditure, and is the most widely used source in studies of consumption requiring panel data. However, recent work by Aguiar and Hurst (2005) show that food expenditure is unlikely to be a good proxy for actual consumption.

[10] See their Table 3, pg. 708.

[11] See their footnote 19, pg 21, which refers to the estimated autocovariances from their Table VI.

[12] They specify a Markov transition matrix with 9 bins (corresponding to 9 quantiles) for consumption. We thank Fabrizio Perri for providing us with their estimated matrix. Using this matrix we computed that the conditional variance of consumption growth had an average across bins of 0.0646 (this is for the year 2000, the last in their sample; but the results are similar for other years).

specifications for consumption, the problem is that the key parameter remains largely unknown. This suggests that a preferable strategy is to use consumption processes obtained from models that have been successful at matching the available data on consumption and income.

Lessons. Two lessons emerge from our simple exercise. First, welfare gains are potentially far from trivial. Second, they are quite sensitive to two parameters available in our exercise: the variance in the growth rate of consumption and the subjective discount factor.

Based on the empirical uncertainty regarding these parameters, the second point suggests obtaining the baseline allocation from a model, instead of attempting to do so directly from the data. In addition, another reason for starting from a model is to address an important caveat in the previous exercise, the assumed linear accumulation technology. As we shall see, with a more standard neoclassical technology welfare gains are generally greatly tempered.

5   General Equilibrium

Up to now the analysis has been that of partial equilibrium. Alternatively one can interpret the results as applying to an economy facing some given constant rate of return to capital. We now argue that this magnifies the welfare gains from reforming the consumption allocation. Specifically, with a concave accumulation technology, as in the neoclassical growth model, the welfare effects may be greatly reduced. This point is certainly not surprising, nor is it specific to the model or forces emphasized here. Indeed, a similar issue arises in the Ramsey literature, the quantitative effects of taxing capital greatly depend on the underlying technology.[13] Unsurprising as it may, it is important to confront this issue to reach meaningful quantitative conclusions.[14]

To make the point clear, consider for a moment the opposite extreme: the economy has no savings technology, so that $C_t \leq N_t$ for $t = 0, 1, \ldots$ (Huggett, 1993). In this case, if the utility function is of the CRRA form then a baseline geometric random walk allocation is constrained efficient. This follows using Proposition 2 since one can verify that equation (9) holds. Thus, in this exchange economy there are no welfare gains from changing the allocation.

---

[13] Indeed, Stokey and Rebelo (1995) discuss the effects of capital taxation in representative agent endogenous growth models. They show that the effects on growth depend critically on a number of model specifications. They then argue in favor of specifications with very small growth effects, suggesting that a neoclassical growth model with exogenous growth may provide an accurate approximation.

[14] A similar point is at the heart of Aiyagari's (1994) paper, which quantified the effects on aggregate savings of uncertainty with incomplete markets. He showed that for given interest rates the effects could be enormous, but that the effects were relatively moderate in the resulting equilibrium of the neoclassical growth model.

Certainly the fixed endowment case is an extreme example, but it serves to illustrate that general equilibrium considerations are extremely important. For a neoclassical growth model the welfare gains lie somewhere between the endowment economy and the economy with a fixed rate of return; where exactly, is precisely what we explore next.

5.1 Log Utility: Separation of Aggregate and Idiosyncratic

Aggregate and Idiosyncratic Subproblems. We now derive a strong separation result that greatly simplifies the general equilibrium problem for the logarithmic utility function case. To simplify the presentation, we shall assume that the baseline allocation features constant aggregate labor supply N_t; this is implied whenever the baseline allocation is at a steady state. To simplify the notation, we drop N and write the resource constraints as $G(K_{t+1}, K_t, C_t) \geq 0$ for $t = 0, 1 \ldots$

First, without loss in generality, one can always decompose any allocation into an idiosyncratic and an aggregate component: $\tilde{u}_t(\theta^t) = \hat{u}_t(\theta^t) + \delta_t$; the aggregate component $\{\delta_t\}$ is a deterministic sequence; the idiosyncratic component $\{\hat{u}_t\}$ has constant expected consumption normalized to one: $\mathbb{E}[c(\hat{u}_t(\theta^t))] = 1$. Since our feasible perturbations certainly allow parallel deterministic shifts it follows that

$$\{\tilde{u}_t\} \in \Upsilon(\{u_t\}, \Delta_{-1}) \quad \Longleftrightarrow \quad \{\hat{u}_t\} \in \Upsilon\Big(\{u_t\}, \Delta_{-1} + \sum_{t=0}^{\infty} \beta^t \delta_t\Big).$$

Since Δ_{-1} is a free variable in the General Equilibrium Problem, this implies that we only need to ensure that $\{\hat{u}_t\}$ is a feasible perturbation for some value of the free variable $\hat{\Delta}_{-1} = \Delta_{-1} + \sum_{t=0}^{\infty} \beta^t \delta_t$.

All the arguments up to this point apply for any utility function. However, with logarithmic utility $\delta_t = \log(C_t)$ and it follows that the planning problem can be decomposed as follows.

General Equilibrium Problem with Log utility

$$\max_{\{\tilde{u}_t\},C_t,K_{t+1},\hat{\Delta}_{-1}} \left[\sum_t \beta^t \mathbb{E}[\hat{u}_t(\theta^t)] + \sum_t \beta^t U(C_t) \right]$$

subject to,

$$G(K_{t+1}, K_t, C_t) \geq 0 \qquad t = 0, 1, \ldots \tag{22}$$

$$\mathbb{E}[c(\hat{u}_t(\theta^t))] = 1 \qquad t = 0, 1, \ldots \tag{23}$$

$$\{\hat{u}_t\} \in \Upsilon(\{u_t\}, \hat{\Delta}_{-1})$$

It is now apparent that we can study the idiosyncratic component problem of choosing $\{\hat{u}_t\}$, from the aggregate one of of selecting $\{C_t, K_{t+1}\}$. The aggregate variables $\{C_t, K_{t+1}\}$ maximize the objective function subject only to the resource constraint (22).

Aggregate Component Problem

$$\max_{C_t, K_{t+1}} \sum_t \beta^t U(C_t)$$

subject to

$$G(K_{t+1}, K_t, C_t) \geq 0 \qquad t = 0, 1, \ldots$$

In other words, the problem is simply that of a standard deterministic growth model, which, needless to say, is straightforward to solve.

The idiosyncratic component problem finds the best perturbation with constant consumption over time.

$$\max_{\{\tilde{u}_t\},\hat{\Delta}_{-1}} \sum_t \beta^t \mathbb{E}[\hat{u}_t(\theta^t)]$$

subject to

$$\mathbb{E}[c(\hat{u}_t(\theta^t))] = 1 \qquad t = 0, 1, \dots$$
$$\{\hat{u}_t\} \in \Upsilon(\{u_t\}, \hat{\Delta}_{-1})$$

Moreover, the idiosyncratic problem can be solved by solving a partial equilibrium problem with with $q = \beta$. This follows since the Inverse Euler equation then implies that $\mathbb{E}_t[\tilde{c}_{t+1}] = c_t$, so that aggregate consumption is constant. Hence, we can find the solution using Proposition 4.

5.2   Quantifying the Welfare Gains

In this section we explore welfare gains quantitatively taking general equilibrium effects into full account, using the methodology developed above for logarithmic utility. We first replicate Aiyagari's (1994) seminal incomplete markets exercise. We then take the general equilibrium allocation from this model as our baseline.

Aiyagari considered a Bewley economy—where a continuum of agents each solve an income fluctuations problem—imbedded within the neoclassical growth model. The time horizon is infinite. There is no aggregate uncertainty, yet individuals are confronted with significant idiosyncratic risk: after-tax labor income is stochastic. Efficiency labor is specified directly as a first-order autoregressive process in logarithms:

$$\log(n_t) = \rho \log(n_{t-1}) + (1 - \rho) \log(\bar{n}) + \varepsilon_t$$

where ε_t is an i.i.d. random variable assumed Normally distributed. With a continuum of agents the average efficiency labor supplied is $\bar{n}$. Labor income is given by the product $w \, n_t$ where w is the steady-state wage.

There are no market insurance arrangements, so agents must cope with their risk. They can do so by accumulating assets, and possibly borrowing, at a constant interest rate r. The budget constraints are

$$a_{t+1} + c_t \le (1 + r)a_t + w n_t$$

23

for all $t = 0, 1, \ldots$ and histories of shocks. In addition the agent is subject to some borrowing limit $a_t \geq \underline{a}$; we take Aiyagari benchmark, where there is no borrowing $\underline{a} = 0$.

The equilibrium steady-state wage is given by the marginal product of labor $w = F_n(K_{ss}, \bar{n})$ and the interest rate is given by the net marginal product of capital, $r = F_K(K_{ss}, \bar{n}) - \delta$.[15] For any given interest rate, individual saving behavior leads to an invariant cross-sectional distribution of asset holdings. At a steady-state equilibrium the interest rate induces a distribution with average assets equal to the capital stock K. The individual consumption allocation is a function of assets and current income, so that it can be written as a function of the state variable $s_t \equiv (a_t, y_t)$ which evolves as a Markov process.

Table 1 shows the computed equilibrium interest rates and the implied welfare properties with logarithmic utility; the rest of the parameters are quite standard: the discount factor $\beta = .96$ the production function is Cobb-Douglas with a share of capital set to 0.36, capital depreciation is set to 0.08. This is done for various combinations of the autocorrelation coefficient and the standard deviation of income growth (comparable to Aiyagari's Table II, pg. 678). We also break down the total welfare gains into the the idiosyncratic and aggregate components.

Aiyagari argues, based on various sources of empirical evidence, for a parameterization with a coefficient of autocorrelation of $\rho = 0.6$ and a standard deviation of labor income growth of 20%. For this preferred specification, we find that welfare gains are minuscule. This contrasts sharply with the partial equilibrium exercises and illustrates the importance of general equilibrium effects. Relatively small welfare gains are also obtained for higher values of ρ or the variance of income growth. Welfare gains become non-negligible, up to around 1%, only when income shocks display extreme persistence and variance.

Overall, welfare gains appear to be very modest—especially when compared to the partial equilibrium exercise in the previous section. To understand these findings it is useful to discuss the idiosyncratic and aggregate components separately. Our finding that idiosyncratic gains are modest could have perhaps been anticipated by our illustrative geometric random-walk example, where idiosyncratic gains are zero. Intuitively, welfare gains from the idiosyncratic component require differences in the expected consumption growth rate across individuals. When individuals smooth their consumption over time effectively the remaining differences are small—as a result, so are the welfare gains.

Welfare gains from the aggregate component are directly related to the difference between the equilibrium and optimal stead-state capital. With logarithmic utility this is equivalent,

---

[15] For simplicity this assumes no taxation. It is straightforward to introduce taxation. However, we conjecture that since taxation of labor income acts as insurance, it effectively reduces the variance of shocks to net income. Lower uncertainty will then only lower the welfare gains we compute.

to the difference between the equilibrium steady-state interest rate and $\beta^{-1} - 1$, the interest rate that obtains with complete markets. Hence, our finding of low aggregate welfare gains is directly related to Aiyagari's (1994) main conclusion: precautionary savings are small in the aggregate, in that steady-state capital and interest rate are close their complete-markets levels, as shown in our Table 1. Our exercise establishes that general equilibrium effects crucial, not just for such positive quantitative conclusions, but also for normative ones.

| | | Std($\log(n_t) - \log(n_{t-1})$) = .2 | | |
| | | | Welfare Gains | |
ρ	r	Idiosyncratic	Aggregate	Total
0	4.1467%	.006%	$\sim 0\%$	0.006%
0.3	4.1260%	0.018%	$\sim 0\%$	0.018%
0.6	4.0856%	0.05%	0.001%	0.051%
0.9	3.9493%	0.19%	0.008%	0.198%

| | | Std($\log(n_t) - \log(n_{t-1})$) = .4 | | |
| | | | Welfare Gains | |
ρ	r	Idiosyncratic	Aggregate	Total
0	4.0149%	0.06%	0.004%	0.064%
0.3	3.9067%	0.15%	0.012%	0.162%
0.6	3.7113%	0.40%	0.039%	0.439%
0.9	3.3036%	1.16%	0.15%	1.31%

Table 1: Welfare Gains for replication of Aiyagari (1994).

To get a feel for the magnitudes, Figure 3 computes aggregate welfare gains as a function of the initial condition, expressed in terms of the difference between the baseline steady-state rate of interest $R^{ss} = F'(k_0) + 1 - \delta$ and β^{-1}. The gains are non-linear and convex shaped. The figure illustrates why the gains from the aggregate component are quite modest in Table 1: the largest difference in interest rates is less than 1% in the table, but it takes a much larger difference, of around 2%, to get welfare gains that are bigger than 1%.

5.3 General Utility

Relaxed Problem. A fruitful way of attacking the problem for general preferences is to consider a relaxed version that replaces the sequence of constraints $C_t = \mathbb{E}[c(\tilde{u}_t(\theta^t))]$ for $t = 0, 1, \ldots$ with a single present value condition. This *relaxed problem* takes as given a sequence of intertemporal prices $\{Q_t\}$ used to compute the present value.
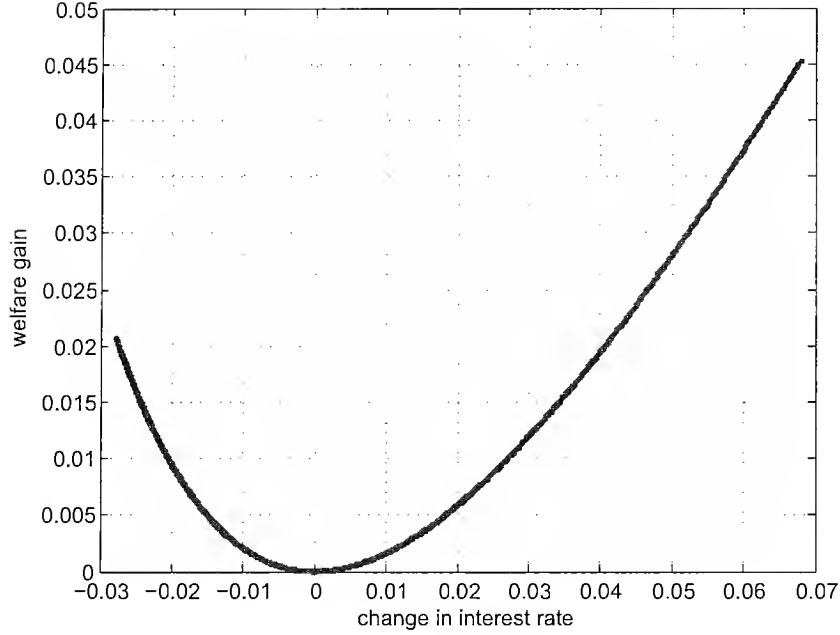
Figure 3: Aggregate component of welfare gains for logarithmic utility as a function of the initial interest rate difference: $(F'(k_0) + 1 - \delta) - \beta^{-1}$.

Relaxed Problem

$$\max_{\{\tilde{u}_t\}, C_t, K_{t+1}, \Delta_{-1}} \sum_t \beta^t \mathbb{E}[\tilde{u}_t(\theta^t)]$$

subject to,

$$G(K_{t+1}, K_t, C_t) \geq 0 \qquad t = 0, 1, \ldots$$

$$\sum_{t=0}^{\infty} Q_t C_t = \sum_{t=0}^{\infty} Q_t \mathbb{E}[c(\tilde{u}_t(\theta^t))] \tag{24}$$

$$\{\tilde{u}_t\} \in \Upsilon(\{u_t\}, \Delta_{-1})$$

The connection with the original problem is the following. If the prices $\{Q_t\}$ induce a solution to the relaxed problem that has $C_t = \mathbb{E}[c(\tilde{u}_t(\theta^t))]$ holding for all $t = 0, 1, \ldots$ then this also solves the original problem. Indeed, one can interpret $\{Q_t\}$ as Lagrange multipliers; appealing to a Lagrangian necessity theorem then ensures the existence of such a sequence $\{Q_t\}$. This 'relaxed problem' approach is adopted from Farhi and Werning (2005).

The vantage point of this approach is that it decouples the aggregate and idiosyncratic

26

parts of the problem. As a result, as long as the baseline allocation can be written recursively as a function of some state variable s, one can solve for the consumption allocation using a Bellman equation similar to the one from Section 2. To see this, note that the relaxed problem must solve the dual that minimizes expected discounted costs, given by (24), subject to delivering a certain promised lifetime utility level. If one parameterizes the promised utility level, relative to the baseline, by Δ_{-1} and let $q_t \equiv Q_{t+1}/Q_t$ we obtain the following representation.

Nonstationary Bellman equation

$$K(\Delta; s, t) = \min_{\Delta'} \left[c(u(s) + \Delta - \beta\Delta') + q_t \mathbb{E}[K(\Delta'; s', t+1) \mid s] \right]$$

Indeed, an economy that converges to a steady-state has q_t converging to some constant value q, so a stationary Bellman as in Section 2 then characterizes the long run. In this sense, the analysis of the dynamic programming problem with given q remains relevant.

As for the aggregate sequence of capital it is trivially determined by the first order condition from the relaxed problem. This directly pins down the sequence $\{K_t\}$ given the sequence of prices $\{Q_t\}$. Reversely, any sequence of capital $\{K_t\}$ can be associated one-to-one with a sequence of relative prices $\{Q_{t+1}/Q_t\}$.

In principle, the entire problem, including the transitional dynamics, can then be solved numerically by a shooting algorithm taking aim at the identified steady state, which can be found using the method behind Figure 1. One can compute simple upper and lower bounds on the welfare gains that include the transition, without the need to fully solve for it.

Simple Upper and Lower Welfare Bounds

Simple upper and lower bounds can be computed for the welfare gains that are likely to be very informative. We suppose the baseline allocation initially finds itself at a steady state.

Upper Bound. We start with the upper bound. The idea is simple, we replace the actual concave production function $F(k) + (1 - \delta)k$ with a linear one tangent at the new steady state: $\bar{F}(k) \equiv F'(K_{ss})(K - K_{ss}) + F(K_{ss}) + (1 - \delta)K_{ss}$. We then solve the planning problem given this technology. This problem is simple because it is equivalent to a stationary relaxed problem with $q = F'(k_{ss})^{-1}$. The welfare improvement then constitutes an upper bound to the true gains since the technology $\bar{F}$ is better than F.

Lower Bound. Moving on to the lower bound, the idea is simply to not solve the full maximization. Instead, one can take any allocation for utility $\{\hat{u}_t\} \in \Upsilon(\{u_t\}, \Delta_{-1})$, and

27

for any such sequence consider the best feasible allocation generated from parallel shifts $\tilde{u}_t(\theta^t) = \hat{u}_t(\theta^t) + v_t$ for a deterministic sequence $\{v_t\}$. Define the sequence of consumption functions $\Psi_t(v_t) = \mathbb{E}[c(\hat{u}_t(\theta^t) + v_t)]$. Then the problem reduces to the deterministic general equilibrium problem

$$\max_{\{C_t\}} \sum_{t=0}^{\infty} \beta^t \Psi_t^{-1}(C_t)$$

$$G(K_{t+1}, K_t, C_t) \geq 0 \qquad t = 0, 1, \ldots$$

and welfare is the value of this problem and the discounted expected value from the allocation $\{\hat{u}_t\}$. The allocation $\{\hat{u}_t\}$ used in this exercise can be any feasible perturbation, but there are sensible and simple choices. For example, one can simply use the baseline allocation, or alternatively, the solution from the relaxed problem for some arbitrary sequence of prices $\{q_t\}$. A simple case would be to use a constant value of q, say, that from the new steady state q_{ss}. One can indeed use many allocations to produce many such lower bounds and use the highest value thus obtained.

6 Overlapping Generations

We now extend the analysis to overlapping generation frameworks, that can potentially capture rich life-cycle elements, missed by infinitely lived agent models.[16] We use a very general setup that can accommodate many demographical specifications. For example, it nests the canonical two-period structure introduced by Samuelson (1958) and Diamond (1965), Blanchard's (1985) parsimonious perpetual-youth model, or calibration exercises using empirical survival probabilities.

Demographics Agents of generation s have a probability $\lambda_{s,t} \geq 0$ of dying at date t conditional on being alive at date $t-1$, where $\lambda_{s,t} = 0$ for $t \leq s$.[17] Define the survival probability $\pi_{s,t}$ for an agent of generation s at date t:

$$\pi_{s,t} = \prod_{u=s}^{t} (1 - \lambda_{s,u})$$

---

[16] There is a lot of recent work calibrating and estimating overlapping generation models. Gourinchas and Parker (2002) estimate a life-cycle model. Storesletten, Telmer and Yaron (2004) and Heathcote, Storesletten and Violante (2004a) calibrate overlapping generations to address the effects of the rise in inequality in the US during the 80s.

[17] The Samuelson-Diamond two-period specification has $\lambda_{s,s+1} = 0$ and $\lambda_{s,s+2} = 1$. The Blanchard specification assumes $\lambda_{s,t} = (1-p)^{t-s-1}p$ for all $t \geq s+1$.

if $s \leq t$ and $\pi_{s,t} = 0$ if $s > t$. Note in particular that $\pi_{s,s} = 1$. Let Λ_s be the initial size of generation s.[18]

Planning Problem. For any generation s, define $\Gamma_s = \{t \mid t \geq \max(s, 0)\}$, as the future dates where generation s has some live members. Let the relevant history of shocks for an agent in generation s be denote as $\theta_s^t = (\theta_{s,s}, \theta_{s,s+1}, ..., \theta_{s,t})$. The baseline allocation is then a sequence of utility assignments as a function of the history $\{u_{s,t}(\theta_s^t)\}_{t \in \Gamma_s}$ for all generations s. We can accommodate heterogeneity at birth or at the initial date, by allowing a non trivial initial distribution of $\theta_{s,s}$.

We allow for a family of social welfare functions

$$W \equiv \sum_{s=-\infty}^{\infty} \alpha_s W_s \Lambda_s$$

where

$$W_s \equiv \sum_{t \in \Gamma_s} \pi_{s,t} \beta^{t-s} \mathbb{E}[u_{s,t}(\theta_s^t)]$$

The problem of the planner is to maximize the welfare function W subject to

$$\{\tilde{u}_{s,t}\}_{t \in \Gamma_s} \in \Upsilon(\{u_{s,t}\}_{t \in \Gamma_s}, \Delta_s) \qquad \forall s$$

$$G\left(K_{t+1}, K_t, \sum_{s=-\infty}^{t} \Lambda_s \pi_{s,t} C_{s,t}\right) \geq 0 \qquad \forall t$$

$$C_{s,t} = \mathbb{E}[c(\tilde{u}_{s,t})] \qquad \forall t \in \Gamma_s \quad \forall s$$

This problem maximizes the social welfare function within the set of feasible perturbations subject to an aggregate resource constraint. The variable Δ_s represents the change in welfare for generation s. Hence the objective function is equivalent up to a constant (representing welfare at the baseline) to

$$\sum_{s=-\infty}^{\infty} \alpha_s \Lambda_s \Delta_s$$

Note that, as in Section 5, here the initial perturbations Δ_s are free variables in the maximization.

Separation of Aggregate and Idiosyncratic

As in the infinite lived agent case, we can derive separation results that simplify the planning problem. For the sake of brevity, we only discuss the stronger separation result that is

---

[18] For steady states it is natural to suppose a constant population size, $\sum_{s=-\infty}^{\infty} \Lambda_s \lambda_{s,t} = \Lambda_t$.

possible with logarithmic utility.

First, without loss in generality, one can always decompose any allocation into an idiosyncratic and an aggregate component: $u_{s,t} = \hat{u}_{s,t} + \delta_{s,t}$; the aggregate component $\{\delta_{s,t}\}$ is a deterministic sequence; the idiosyncratic component $\{\hat{u}_{s,t}\}$ has constant expected consumption normalized to one: $\mathbb{E}[c(\hat{u}_{s,t})] = 1$ for all s and all t. Since our feasible perturbations certainly allow parallel deterministic shifts it follows that

$$\{\tilde{u}_{s,t}\}_{t \in \Gamma_s} \in \Upsilon(\{u_{s,t}\}_{t \in \Gamma_s}, \Delta_s) \quad \Leftrightarrow \quad \{\tilde{u}_{s,t}\}_{t \in \Gamma_s} \in \Upsilon\left(\{u_{s,t}\}_{t \in \Gamma_s}, \Delta_s + \sum_{t=s}^{\infty} \pi_{s,t} \beta^{t-s} \delta_{s,t}\right)$$

Since Δ_s is a free variable in the General Equilibrium Problem, this implies that we only need to ensure that $\{\hat{u}_{s,t}\}$ is a feasible perturbation for some value of the free variable $\Delta_s + \sum_{t=0}^{\infty} \pi_{s,t} \beta^{t-s} \delta_{s,t}$. In the logarithmic case $\delta_{s,t} = \log(C_{s,t})$ and the original problem can be decomposed as follows.

General Equilibrium Problem with Log Utility. The problem is to maximize

$$\sum_{s=-\infty}^{\infty} \alpha_s \sum_{t=0}^{\infty} \left(\Lambda_s \pi_{s,t} \beta^{t-s} \mathbb{E}[\hat{u}_{s,t}] + \Lambda_s \pi_{s,t} \beta^{t-s} \log(C_{s,t})\right)$$

subject to,

$$\{\hat{u}_{s,t}\}_{t \in \Gamma_s} \in \Upsilon(\{u_{s,t}\}_{t \in \Gamma_s}, \hat{\Delta}_s) \qquad \forall s$$

$$G\left(K_{t+1}, K_t, \sum_{s=-\infty}^{t} \Lambda_s \pi_{s,t} C_{s,t}\right) \geq 0 \qquad \forall t$$

$$1 = \mathbb{E}[c(\hat{u}_{s,t})] \qquad \forall t \in \Gamma_s \quad \forall s$$

where the maximization is performed over $\{\tilde{u}_{s,t}\}, C_{s,t}, K_{t+1}, \hat{\Delta}_{s,s-1}$ for all $s = \ldots, -1, 0, 1, \ldots$ and $t = 0, 1, \ldots$

It is now apparent that we can separate the idiosyncratic and aggregate components. We define the idiosyncratic component problem as seeking finds the best perturbation with constant consumption over time for each generation. It is trivial to see that it can be decomposed into a series of subproblem, one for each generation s.

Idiosyncratic Component Problem

$$\max_{\{\tilde{u}_{s,t}\},\hat{\Delta}_s} \sum_{t=s}^{\infty} \pi_{s,t}\beta^{t-s}\mathbb{E}[\hat{u}_{s,t}]$$

subject to,

$$\{\hat{u}_{s,t}\}_{t\in\Gamma_s} \in \Upsilon(\{u_{s,t}\}_{t\in\Gamma_s}, \hat{\Delta}_s)$$

$$1 = \mathbb{E}[c(\hat{u}_{s,t})]$$

The idiosyncratic problem can be solved recursively in much the same way as in the infinitely lived agent case.

Aggregate Component Problem. The planner maximizes

$$\max_{C_{s,t},K_{t+1}} \sum_{s=-\infty}^{\infty} \alpha_s \sum_{t=0}^{\infty} \Lambda_s \pi_{s,t}\beta^{t-s}\log(C_{s,t})$$

subject to

$$G\left(K_{t+1}, K_t, \sum_{s=-\infty}^{t} \Lambda_s \pi_{s,t} C_{s,t}\right) \geq 0 \qquad t = 0, 1, \dots$$

where the maximization is performed over $\{C_{s,t}, K_{t+1}\}$.

This problem is a simple deterministic planning problem. Note that its solution will generally depend on the Pareto weights $\{\alpha_s\}$ used in the welfare criterion. An interesting special case is when these weights are such that the baseline sequence of capital K_{t+1} and aggregate consumption $\bar{C}_t \equiv \sum_{s=-\infty}^{t} \Lambda_s \pi_{s,t} C_{s,t}$ are optimal. Welfare gains are then entirely from intra-period allocation problem: for each t distributing $\bar{C}_t$ across generations, choosing $C_{s,t}$ for all s. The necessary weights and the associated welfare gains can actually be solved in closed form in this case.

Appendix

A   Proof of Proposition 4

With logarithmic utility the Bellman equation is

$$K(s, \Delta) = \min_{\Delta'}[s \exp(\Delta - \beta\Delta') + q\mathbb{E}[K(s', \Delta') \mid s]]$$
$$= \min_{\Delta'}[s \exp((1 - \beta)\Delta + \beta(\Delta - \Delta')) + q\mathbb{E}[K(s', \Delta') \mid s]]$$

Substituting that $K(\Delta, s) = k(s) \exp((1 - \beta)\Delta)$ gives

$$k(s) \exp((1 - \beta)\Delta) = \min_{\Delta'}[s \exp((1 - \beta)\Delta + \beta(\Delta - \Delta')) + q\mathbb{E}[k(s') \exp((1 - \beta)\Delta') \mid s]],$$

and cancelling terms:

$$k(s) = \min_{\Delta'}[s \exp(\beta(\Delta - \Delta')) + q\mathbb{E}[k(s') \exp((1 - \beta)(\Delta' - \Delta)) \mid s]]$$
$$= \min_{d}[s \exp(-\beta d) + q\mathbb{E}[k(s') \exp((1 - \beta)d) \mid s]]$$
$$= \min_{d}[s \exp(-\beta d) + q\mathbb{E}[k(s') \mid s] \exp((1 - \beta)d)]$$

where $d \equiv \Delta' - \Delta$. We can simplify this one dimensional Bellman equation further. Define $\hat{q}(s) \equiv q\mathbb{E}[k(s') \mid s]/s$ and

$$M(\hat{q}) \equiv \min[\exp(-\beta d) + \hat{q} \exp((1 - \beta)d)].$$

The first-order conditions gives

$$\beta \exp(-\beta d) = \hat{q}(1 - \beta) \exp((1 - \beta)d) \quad \Rightarrow \quad d = \log \frac{\beta}{(1 - \beta)\hat{q}}. \tag{25}$$

Substituting back into the objective we find that

$$M(\hat{q}) = \frac{1}{1 - \beta} \exp(-\beta d) = \frac{1}{1 - \beta} \exp\left(-\beta \log \frac{\beta}{(1 - \beta)\hat{q}}\right)$$
$$= \frac{1}{(1 - \beta)^{1-\beta}\beta^\beta} \hat{q}^\beta = B\hat{q}^\beta,$$

where B is a constant defined in the obvious way in terms of β.

The operator associated with the Bellman equation is then

$$T[k](s) = sM\left(q\frac{\mathbb{E}[k(s') \mid s]}{s}\right) = As^{1-\beta}\left(\mathbb{E}[k(s') \mid s]\right)^{\beta},$$

where $A \equiv Bq^{\beta} = (q/\beta)^{\beta}/(1-\beta)^{1-\beta}$.

Combining the Bellman $k(s)/s = M(\hat{q}) = A\hat{q}^{\beta}$ with equation (25) yields the policy function as a function of $K(s)$. This completes the proof.

References

Aguiar, Mark and Erik Hurst, "Consumption versus Expenditure," *Journal of Political Economy*, 2005, *113* (5), 919–948. 19

Aiyagari, S. Rao, "Uninsured Idiosyncratic Risk and Aggregate Saving," *Quarterly Journal of Economics*, 1994, *109* (3), 659–684. 1, 3, 10, 11, 12, 13, 20, 23, 24, 25

Albanesi, Stefania and Christoper Sleet, "Dynamic Optimal Taxation with Private Information," 2004. mimeo. 2

Blanchard, Olivier J, "Debt, Deficits, and Finite Horizons," *Journal of Political Economy*, 1985, *93* (2), 223–47. 28

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption inequality and partial insurance," IFS Working Papers W04/28, Institute for Fiscal Studies 2004. 19

Chamley, Christophe, "Optimal Taxation of Capital in General Equilibrium," *Econometrica*, 1986, *54*, 607–22. 2

Conesa, Juan Carlos and Dirk Krueger, "On the Optimal Progressivity of the Income Tax Code," *Journal of Monetary Economics*, 2005, *forthcoming*. 3

Davila, Julio, Jay H. Hong, Per Krusell, and Jose-Victor Rios-Rull, "Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks," PIER Working Paper Archive 05-023, Penn Institute for Economic Research, Department of Economics, University of Pennsylvania July 2005. available at http://ideas.repec.org/p/pen/papers/05-023.html. 3

Deaton, Angus and Christina Paxson, "Intertemporal Choice and Inequality," *Journal of Political Economy*, 1994, *102* (3), 437–67. 19

Diamond, Peter A., "National Debt in a Neoclassical Growth Model," *American Economic Review*, 1965, *55*, 1126–50. 28

— **and James A. Mirrlees**, "A Model of Social Insurance With Variable Retirement," Working papers 210, Massachusetts Institute of Technology, Department of Economics 1977. 2, 4

Farhi, Emmanuel and Iván Werning, "Inequality, Social Discounting and Progressive Estate Taxation," NBER Working Papers 11408 2005. 26

Geanakoplos, John and Heracles M. Polemarchakis, "Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When the Asset Market Is Incomplete," Cowles Foundation Discussion Papers 764, Cowles Foundation, Yale University 1985. available at http://ideas.repec.org/p/cwl/cwldpp/764.html. 3

Golosov, Mikhail, Aleh Tsyvinski, and Iván Werning, "New Dynamic Public Finance: A User's Guide," *forthcoming in NBER Macroeconomics Annual 2006*, 2006. 3

_ , Narayana Kocherlakota, and Aleh Tsyvinski, "Optimal Indirect and Capital Taxation," *Review of Economic Studies*, 2003, *70* (3), 569–587. 2, 4, 6, 8

Gourinchas, Pierre-Olivier and Jonathan A. Parker, "Consumption Over the Life Cycle," *Econometrica*, 2002, *70* (1), 47–89. 28

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante, "The Cross-Sectional Implications of Rising Wage Inequality in the United States," CEPR Discussion Papers 4296, C.E.P.R. Discussion Papers 2004. 28

_ , _ , and _ , "Two Views of Inequality Over the Life-Cycle," *Journal of the European Economic Association (Papers and Proceedings)*, 2004, *3*, 543–52. 19

Huggett, Mark, "The risk-free rate in heterogeneous-agent incomplete-insurance economies," *Journal of Economic Dynamics and Control*, 1993, *17* (5–6), 953–969. 12, 20

Judd, Kenneth L., "Redistributive Taxation in a Perfect Foresight Model," *Journal of Public Economics*, 1985, *28*, 59–83. 2

Kocherlakota, Narayana, "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation," 2004. mimeo. 2

Krueger, Dirk and Fabrizio Perri, "On the Welfare Consequences of the Increase in Inequality in the United States," in Mark Gertler and Kenneth Rogoff, eds., *NBER Macroeconomics Annual 2003*, Cambridge, MA: MIT Press, 2004, pp. 83–121. 19

Ligon, Ethan, "Risk Sharing and Information in Village Economics," *Review of Economic Studies*, 1998, *65* (4), 847–64. 2, 9

Rebelo, Sergio, "Long-Run Policy Analysis and Long-Run Growth," *Journal of Political Economy*, 1991, *99* (3), 500–521. 5

Rogerson, William P., "Repeated Moral Hazard," *Econometrica*, 1985, *53* (1), 69–76. 2, 9

Samuelson, Paul A., "An Exact Consumption-Loan Model of Interest with or without the Social Contrivance of Money," *Journal of Political Economy*, 1958, *66*, 467–82. 28

Slesnick and Ulker, "Inequality and the Life-Cycle: Consumption," Technical Report 2004. 19

Spear, Stephen E. and Sanjay Srivastava, "On Repeated Moral Hazard with Discounting," *Review of Economic Studies*, 1987, *54* (4), 599–617. 12

Stokey, Nancy L. and Sergio Rebelo, "Growth Effects of Flat-Rate Taxes," *Journal of Political Economy*, 1995, *103* (3), 519–50. 20

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron, "Cyclical Dynamics in Idiosyncratic Labor Market Risk," *Journal of Political Economy*, 2004, *112* (3), 695–717. 15, 19, 28

Werning, Iván, "Optimal Dynamic Taxation," 2002. PhD Dissertation, University of Chicago. 2